

Organization 2600 Bldg/Room

U. S. Department of Commerce

Commissioner for Patents


P.O. Box 1450

Alexandria, VA 22313-1450

If Undeliverable Return in Ten Days

AN EQUAL OPPORTUNITY EMPLOYER

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, \$300



ADRESSE INCOMPLÈTE. VEUILLEZ NOUS PRÉCISER
QUELS SERVICES, LABORATOIRES OU INSTITUTS
CELA CONCERNE, ET SURTOUT LE NUMÉRO DE LA **STATION**

Zutreffendes durchkreuzen - Marquer ce qui convient
Porre una crocetta secondo il caso - Tick off the corresponding space

Abgereist Parti Partito Left	Adresse ungenügend Address insufficiente Indirizzo insufficiente Address Insufficient	Unbekannt Inconnu Sconosciuto Unknown	Annahme verweigert Refusé Respinto Refused	Gestorben Décédé Decesso Deceased
---------------------------------------	--	--	--	--



UNITED STATES PATENT AND TRADEMARK OFFICE

IFW

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

09/902,227

07/11/2001

Roger D. Hersch

7585

7590 01/27/2009
Prof. Roger D. Hersch
EPFL - DI/LSP - INF Ecublens
~~CH-1015 - Lausanne,~~
SWITZERLAND

EXAMINER

ROSARIO, DENNIS

ART UNIT

PAPER NUMBER

2624

MAIL DATE

DELIVERY MODE

01/27/2009

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 09/902,227	Applicant(s) HERSCH ET AL.	
	Examiner Dennis Rosario	Art Unit 2624	

– The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 04 July 2007.
2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1,3-5,7,10,13,24-28 and 34-3 is/are pending in the application.
4a) Of the above claim(s) _____ is/are withdrawn from consideration.
5) ☐ Claim(s) _____ is/are allowed.
6) ☒ Claim(s) 1,3-5,7,10,13,24-28 and 34-38 is/are rejected.
7) ☐ Claim(s) _____ is/are objected to.
8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
10) ☒ The drawing(s) filed on 11 July 2007 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

Response to Amendment

1. The after final amendment was received on 7/4/07. Claims 1,3-5,7,10,13,24-28 and 34-38 are pending.

Allowable Subject Matter

2. The indicated allowability of claims 1,3-5,7,10,13,24-28 and 34-38 is withdrawn in view of the newly discovered reference(s) to Finkelstein et al. (Image Mosaics) in view of Browne et al. (US Patent 6,504,545 B1). Rejections based on the newly cited reference(s) follow.

Claim Rejections - 35 USC § 101

3. 35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

The USPTO "Interim Guidelines for Examination of Patent Applications for Patent Subject Matter Eligibility" (Official Gazette notice of 22 November 2005), Annex IV, reads as follows:

Nonfunctional descriptive material that does not constitute a statutory process, machine, manufacture or composition of matter and should be rejected under 35 U.S.C. Sec. 101. Certain types of descriptive material, such as music, literature, art, photographs and mere arrangements or compilations of facts or data, without any functional interrelationship is not a process, machine, manufacture or composition of matter. USPTO personnel should be prudent in applying the foregoing guidance. Nonfunctional descriptive material may be claimed in combination with other functional descriptive multi-media material on a computer-readable medium to provide the necessary functional and structural interrelationship to satisfy the requirements of 35 U.S.C. Sec. 101. The presence of the claimed nonfunctional descriptive material is not necessarily determinative of nonstatutory subject matter. For example, a computer that recognizes a particular grouping of musical notes read from memory and upon recognizing that particular sequence, causes another defined series of notes to be played, defines a functional interrelationship among that data and the computing processes performed when utilizing that data, and as such is statutory because it implements a statutory process.

4. Claim(s) 24-28 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter as follows. Claim 24 recites an image which

does not impart functionality to a computer or computing device, and is thus considered nonfunctional descriptive material. Such nonfunctional descriptive material, in the absence of a functional interrelationship with a computer, does not constitute a statutory process, machine, manufacture or composition of matter and is thus non-statutory per se. Thus, claims 25-28 are rejected.

5. Claim(s) 1,3-5,7,10,13,24-28 and 34-38 are rejected under 35 U.S.C. 101 as not falling within one of the four statutory categories of invention. Supreme Court precedent¹ and recent Federal Circuit decisions² indicate that a statutory "process" under 35 U.S.C. 101 must (1) be tied to another statutory category (such as a particular apparatus), or (2) transform underlying subject matter (such as an article or material) to a different state or thing. While the instant claim(s) recite a series of steps or acts to be performed, the claim(s) neither transform underlying subject matter nor positively tie to another statutory category that accomplishes the claimed method steps, and therefore do not qualify as a statutory process. For example, claim 1 does not claim a structural limitation in the body of the claims that performs the methods of the body.

Claim Rejections - 35 USC § 103

6. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the

¹ *Diamond v. Diehr*, 450 U.S. 175, 184 (1981); *Parker v. Flook*, 437 U.S. 584, 588 n.9 (1978); *Gottschalk v. Benson*, 409 U.S. 63, 70 (1972); *Cochrane v. Deener*, 94 U.S. 780, 787-88 (1876).

² *In re Bilski*, 88 USPQ2d 1385 (Fed. Cir. 2008).

invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

7. Claims 1,3-5,7,10,13,24-28 and 34-38 are rejected under 35 U.S.C. 103(a) as being unpatentable over Finkelstein et al. (Image Mosaics) in view of Browne et al. (US Patent 6,504,545 B1).

Regarding claim 34, Finkelstein teaches a computing system for electronically displaying a target image with an embedded microstructure evolving over time, said computing system comprising:

a) a server computing system located at one Internet location for rendering said target image from an original image by synthesizing target image instances and comprising a client computing system located at another Internet location receiving a sequence of image instances from the server computing system and displaying said sequence (not taught in Finkelstein),

b) where a time-dependent animation transformation (implied by "tiles may change over time" on the last page: Video mosaics) specifies how said embedded microstructure (one of said tiles) spatially (not taught in Finkelstein) evolves (corresponding to said change over time) over a succession of displayed target image instances (also known as video of said video mosaics),

c) where said embedded microstructure (or tile) comprises visual motive elements selected from a set of text, logo, symbol (or "Global currency" in fig. 8 that shows an image with currency that typically includes symbols that represent a respective country) and ornament,

d) where said visual motive elements evolve (corresponding to said change) spatially independently (Finkelstein does not clearly teach the claimed element that evolves spatially and independently of a subject) of the content of said original image (such as a globe with no texture or no tile of currency);

e) where synthesizing (implied in a video environment) the target image instances (or video) from said original image (said globe with no texture) comprises a halftoning operation (in section 3 and 4.4 of Finkelstein) which adapts intensities (or a color for "Color correction" in 4.4), respectively colors of said visual motive elements to intensities (or the corrected color), respectively colors of said two-dimensional original image.

Finkelstein does not teach the server limitation, spatially evolves and evolve spatially independently of the content of said original image. However, Finkelstein teaches that the method can be practiced in the environments of advertising and commercials, art, transmission and security. In addition, Finkelstein teaches that the tiles or mosaic tiles can be changed as a function of time; however, Finkelstein does not clearly state how they can change.

Browne teaches "Internet based advertising" in col. 1, lines 11-13 and "mosaic tile colouring" in col. 10, lines 29-31 and the limitations of:

a) a server computing system located at one Internet location for rendering said target image from an original image by synthesizing target image instances and comprising a client computing system located at another Internet location receiving a sequence of image instances from the server computing system and displaying said sequence (the whole limitation is well known to one of ordinary skill in the Internet given that Browne teaches said Internet based advertising)

b) where a time-dependent animation transformation (corresponding to "animations...obey...time" in col. 11, lines 41-45) specifies how said embedded microstructure (a tile of said mosaic tile colouring) spatially evolves (two interpretations from col. 9, lines 44-60:1) Browne teaches overlapping shapes or mosaic tiles, which includes a spatial aspect such as a foreground and background, that evolve for the fact that the shapes or tiles progress forward in time and changes via the animation that obeys time. 2) Browne teaches that the overlapping mosaic tiles change in colour as a function of changing a radius and hence corresponding circle over time; thus, Browne teaches a spatial evolution in the context of a circle's size that changes the color of the overlapping mosaic tiles) over a succession of displayed target image instances (so as to form a result of said animation)

c) where said visual motive elements evolve spatially (via the two interpretations) independently of the content of said original image (as shown in fig. 16(a)(d) where the letter "m" is maintained, but changes in colour via the overlapping mosaic tiles).

It would have been obvious at the time the invention was made to one of ordinary skill in the art to modify Finkelstein's halftoning to create mosaic tiles with Browne use of the mosaic tiles for animation, because Browne's animated tiles can be used in advertising to catch a customer's attention more than a static advertising image in order to sell a product.

Claims 24,35,36 and 37 are rejected the same as claim 34. Thus, argument similar to that presented above for claim 34 is equally applicable to claims 24 and 35-37.

Regarding claim 25, Browne teaches the image of claim 24, where the visibility of the embedded microstructure (or "opacity" in fig. 3) is tuned by a mask (or "opacity function" in col. 9, lines 55-58) whose values (or "variable parameter" in col. 9, lines 51-54) represent relative weights ("opacity index" in col. 9, lines 51-54) of said original image (said letter "m" without coloring of mosaic tiles) without embedded microstructure (without said mosaic tiles) and a corresponding image rendered (any one of fig. 16(a)-16(d)) with the embedded microstructure (since fig. 16 shows the colored mosaic tile of said letter "m").

Claims 26 and 27 is rejected the same as claim 25. Thus, argument similar to that presented above for claim 25 is equally applicable to claims 26 and 27.

Claim 28 is rejected the same as claim 37. Thus, argument similar to that presented above for claim 37 is equally applicable to claim 28.

Regarding claim 38, Browne teaches the computing system of claim 36, where the microstructure evolution parameters (all of fig. 3 as applied to the tiles of Finkelstein)

also comprise a warping transformation (fig. 15(a)-(d)) and where the client computing and display system (as implied by said Internet based advertising) also receives from the server computing system as input data a mask (fig. 3:Opacity) whose values represent relative weights of the original color image and of image instances obtained by said halftoning operation (the result of which is a mosaic image in fig. 16(a) which is based upon the halftone inspired mosaic tiles of Finkelstein), the mask defining the position and visibility of the microstructure within the target image (as shown by the progression from fig. 16(a)-16(d)).

Claim 1 is rejected the same as claim 34. Thus, argument similar to that presented above for claim 34 is equally applicable to claim 1 except for the additional limitation as taught by Finkelstein of where said rendering step (or displaying step) comprises a mapping (corresponding to finding a "most similar" image in section 4.3 of Finkelstein) of positions between target image instances (Finkelstein does not teach mapping between positions between target image instances, which are implied frames of the video mosaics) and positions (or "each location" in 4.3) within said original microstructure space (corresponding to a "location within the tile grid" in section 4.3) according to said time-dependent geometric animation transformation and a halftoning of said two-dimensional original image (so that the most similar image, which is halftone-based, can be used for a video mosaic that implies the claimed transformation).

Finkelstein does not teach mapping between target image instances, but teaches that the mosaic tiles can be used in a video environment with a target image where each frame with the target image and tiles are the claimed target image instances.

Browne teaches the video environment with mosaic tiles and the claimed mapping of positions between target image instances via a frame index that can find for example frame 3 which is in a position between frames 2 and 4 of a frame sequence as discussed in col. 12, lines 3-16.

It would have been obvious at the time the invention was made to one of ordinary skill in the art to modify Finkelstein's video mosaics with Browne's video index, because Browne's index can quickly display a desired image instead of playing the whole animation of the desired image.

Regarding claim 3, Browne teaches the method of claim 1, where only a part of said two-dimensional original image (fig. 16(a)) defined by a mask (or said overlapping images) is rendered (displayed) with said embedded microstructure (as shown by the interior of the letter "m").

Regarding claim 4, Finkelstein teaches the method of claim 1, where an additional step enables to specify a set of basic colors ("three color channels" in section 4.4 such as RGB,HSV,Lab) for rendering said target image instances.

Regarding claim 5, Finkelstein teaches the method of claim 4, where said two-dimensional original image (said letter "m") is halftoned by dithering (or color corrected using the inspiration of a "dither-matrix" in said 4.4 that implies the claimed halftoned)

at least one of the basic colors with a dither matrix (as shown in section 3 as "D") embedding the microstructure (resulting in figures 4(a)-(e)).

Regarding claim 7, Finkelstein teaches the method of claim 4, where halftoning is carried out by multicolor dithering (given that Finkelstein uses a color embodiment and a grayscale embodiment as discussed in 4.4) with the defined set of basic colors and with a dither matrix embedding the microstructure.

Regarding claim 10, Finkelstein teaches the method of claim 1, where the evolution of said embedded microstructure over time (which corresponds to the video mosaic that can change tiles over time) comprises a blending between two microstructure shapes (as known to one of ordinary skill in art of art: "blends" in section 2 a concept of which is applied to tiles, which is rectangular in shape).

Regarding claim 13, Finkelstein teaches the method of claim 1, where the embedded microstructure is made more flexible by an additional warping transformation (corresponding to "other... sophisticated rules" in section 4.4, 3rd paragraph includes a distortion of colors which corresponds to the claimed additional warping transformation, such as warping of colors) mapping between a target image space (or displayed image) containing the target image (to be displayed) and an animated dither matrix space (said inspired halftoning tile that can change in the video mosaic: note that the tile does not actually move, but has the capability to move in the video mosaic. If the claimed animated dither matrix space actually displayed motion, then a mapping from an image displaying motion to the claimed target image would overcome the interpretation of claim 13).

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Dennis Rosario whose telephone number is (571) 272-7397. The examiner can normally be reached on 9-5.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Bella can be reached on (571) 272-7778. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Dennis Rosario/
Examiner, Art Unit 2624

/Matthew C Bella/
Supervisory Patent Examiner, Art
Unit 2624

Application/Control Number: 09/902,227
Art Unit: 2624

Page 12

Notice of References Cited	Application/Control No. 09/902,227		Applicant(s)/Patent Under Reexamination HERSCH ET AL.	
	Examiner Dennis Rosario		Art Unit 2624	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-6,504,545	01-2003	Browne et al.	345/473
	B	US-			
	C	US-			
	D	US-			
	E	US-			
	F	US-			
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	Finkelstein et al., Lecture Notes in Computer Science: Electronic Publishing, Artistic Imaging, and Digital Typography: Image Mosaics, March 30-April 3 1998, Springer Berlin, Vol. 1375/1998, pp. 11-22.
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Image Mosaics

Adam Finkelstein and Marisa Range

Computer Science Department, Princeton University
Princeton, NJ 08540, USA
<http://www.cs.princeton.edu/gfx/proj/mosaic>

Abstract. We describe a process for creating an *image mosaic*—a collection of small images arranged in such a way that when they are seen together from a distance they suggest a larger image. To visually suggest the larger form, the small images are arranged to match a large picture as much as possible, and then their colors are adjusted to better suggest the overall form. Arrangement of the small images may be either manual or automatic. Adjustment of the colors in the small image to further suggest the larger picture is fully automatic and employs a new color correction scheme that generalizes traditional halftoning.

1 Introduction

Painters of the impressionist movement exploited a property of the human visual system that combines colors in a region such that the observer sees an overall average color for that region. When viewed up close, an impressionist painting appears to be a collection of small brush strokes of various colors, whereas at a distance those brush strokes combine to yield an overall impression that is typically the subject of the painting. More recently, artists and photographers have exploited the same principle to produce *image mosaics*—layered imagery, where the subject of the work is both the tiny features that only can be seen up close and the large scale features that only can be seen at a distance. This paper explores the use of computers to automatically or semi-automatically produce such imagery.

In this paper we describe methods for arranging a set of small images that we call *tile images* and adjusting their colors so that together they suggest a larger form, as shown in Figures 1, 2, and 3. The motivation for this work is primarily for artistic purposes. The ability to combine groups of pictures in this way affords opportunities for both aesthetic and associational juxtaposition of images. Additionally, there are other potential applications for this kind of technology that we have not emphasized in this work: encoding extra information in an image for transmission or security, new forms of halftoning screens for printing, and association of images for advertising.

The rest of this paper is organized as follows. In Sections 2 and 3, we describe related work and provide some background on traditional halftoning. Section 4 presents our technique for creating image mosaics. In section 5, we describe the individual images that appear in the paper. Section 6 concludes with areas for future work.




Fig. 1. Mona Lisa (three layers)



Fig. 2. John F. Kennedy

2 Related work

Artists have known for some time how to create pictures out of other pictures. Indeed, the efforts described in this paper were largely inspired by Salvador Dali's lithograph *Lincoln in Dalivision*. In his etching [2], Dali shows his wife Gala looking up through a block window at the Godhead. Up close, Gala is surrounded by a jumble of colors, textures, and small caricatures. From a distance, the entire work blends into a bust of Abraham Lincoln. A similar layered quality is exhibited in *Self Portrait I* by Chuck Close [1] in which the large-scale figure is his own face, but it is composed of hundreds of tiny abstract figures. To produce these works, the artists had to "see" the large-scale image before it was actually created, and then exactly reproduce it by arranging and adjusting the smaller figures from which it was composed. In this paper, we describe a process for creating such images more automatically, while still providing the means for artistic expression through composition.

Several researchers have investigated the use of computers to produce pictures in the style of impressionist paintings. Haeberli described a method of creating an image whose overall form matches a given picture, but is composed of tiny brush strokes [5]. Meier extended this work to apply to 3D animation [9], and Litwinowicz subsequently extended it to video [7]. Like these projects, our work begins with the larger form; however, rather than using small brush strokes, we use small images to convey the larger form.




Fig. 3. *American Gothic* composed of pictures from the Web. (Color image)

The process we describe is similar to the artistic screening technique developed by Ostromoukhov and Hersch [11]. In their approach, variations in brightness across the larger image are produced by varying the sizes and shapes of tiny subjects (for example fish, birds, or abstract blobs) or characters (for example, Roman letters, Kanji, or Islamic calligraphy). The tiny figures are described by closed contours separating black from white, so the entire work is ultimately composed of black and white. In contrast, our method varies the brightness of tile images composed of shades of gray (*grayscale*), so that the resulting mosaic is itself a grayscale image. Because we are working with images, we are also able to show how these techniques generalize from grayscale to color.


Photographers, and more recently researchers in computer vision, have addressed the problem of constructing a *panoramic mosaic*—a single, consistent view of a scene

pieced together from a series of photographs capturing the scene from different perspectives [12]: The objective of panoramic mosaicing is to produce a mosaic which does *not* reveal that it was stitched together from different images. This contrasts with our goal of constructing a mosaic wherein the visual artifacts of the tile images are part of the subject of the work.

Our technique appears to be most similar to that of Silvers [13] who has applied this technology for both artistic and commercial purposes. It is difficult to describe the work precisely, as his techniques remain proprietary. It appears that Silvers focuses most of his effort on finding a suitable arrangement for the tile images (the subject of Section 4.3) and avoids correcting the tile images after arranging them (Section 4.4). Silvers tends to use many more image tiles in his mosaics, and is able to produce stunning reproductions of an original image without altering the tile images.

3 Background

Traditional halftoning employs black dots of varying sizes arranged in a regular grid to convey various shades of gray. Historically, the sizes of the dots have been “calculated” by a purely mechanical process called *screening* the target image is photographed slightly out of focus through a mesh or screen on high-contrast film. With the advent of the digital age, screening is now performed almost exclusively by computers. While there are a host of schemes for computing the sizes, shapes and arrangement of dots of ink that faithfully reproduce the target image, many of them share a common feature called a *dither matrix*. Since this matrix plays a role similar to that of our small images tiles, we’ll describe how it works.



Suppose we have (as shown above) a target image I of width w and height h whose pixels are grayscale values ranging from 0 for black to 1 for white. We wish to produce a $w \times h$ halftoned image H , each pixel $H[x,y]$ of which will either be black (ink) or white (paper). The most common approaches employ a $m \times n$ dithering matrix D , where $m \ll w$ and $n \ll h$. For each pixel $I[x,y]$, we find the entry $D[i,j]$ using:

$$i = x \text{ modulo } m$$

$$j = y \text{ modulo } n$$

We choose black or white dots for our halftoned image as follows: if $I[x,y] < D[i,j]$ then $H[x,y]$ is black; otherwise it is white. The key to designing an effective screening




Fig. 4. (a) original grayscale ramp, (b) clustered-dot dither,
(c-e) image mosaics composed of Lincoln, Mona Lisa, and an eye.

method is finding a good dither matrix D . There is a rich literature on the subject [14] and we will not review it further here. Figure 4a shows a ramp, which is halftoned in Figure 4b using a 16×8 clustered-dot ordered dither [14].

4 Creating image mosaics

In this section we describe how to create an image mosaic. Specifically, we'll address the problem: given a collection of tile images T_i and a target image I , create an image mosaic M that resembles I at a coarse scale but is composed of a tiling of the T_i 's. We solve the problem in four steps: (1) choose images, (2) choose a tiling grid, (3) find an arrangement for the image tiles within the grid, and (4) correct the tiles to match the target image. The first two stages (and often the third) are performed manually, and are typically iterative. These four steps are discussed in the following sections.

4.1 Choosing images

The choice of subject matter is purely an artistic one. Clearly, one of the attractions of image mosaics is that they juxtapose images at different scales. This property may be interpreted through composition of images themselves. For example, in Figure 5, a large




Fig. 5. An abalone shell screened through smaller pictures of abalone shells taken by a scanning electron microscope. (Color image)

picture of an abalone shell is screened through a group of micron-scale photographs of abalone shell from a scanning electron microscope. Alternately, the juxtaposition may be one of association, as in Figure 2 showing John F. Kennedy composed of smaller pictures of Marilyn Monroe. This association may be put to use for commercial purposes, such as advertisement; Silvers [13] created a picture of George Washington using hundreds of tiny credit cards, as part of a campaign for Mastercard.

Some images simply work better than others. Recognition plays an important role, so iconic figures such as political leaders, actors, famous works of art, and well-known scenes are desirable. Furthermore, figures that are easy to recognize at very low resolution tend to work well; surely this was part of Dali's motivation for selecting Lincoln as his subject in his lithograph *Lincoln in Dalivision*. The distribution of colors in an image can also be a factor. Tile images with relatively uniform distributions of brightness tend to be easier to identify when their colors are adjusted as described in Section 4.4.

4.2 Finding a tiling pattern

To find an appropriate arrangement for the image tiles, the first question that needs to be addressed is what kind of tiling will be used for the mosaic. We have investigated only semiregular rectangular grids for simplicity, although other tilings are possible. Hexagonal grids have also been used for traditional halftoning [14]. A more general (non-periodic) tiling might be found that achieves a better match for a given target image. However, optimizing over all possible tilings and all possible arrangements of

small image within that tiling such that they best resemble a target image is computationally expensive and remains a challenge for future work. Furthermore, since the image tiles themselves tend to be rectangles, a rectangular tiling is a natural choice. Arranged in a rectangular grid, all the small image tiles must have the *same* aspect ratio. Thus, it is important to choose an aspect ratio appropriate for the entire set of image tiles, as they will either have to be cropped or stretched to conform to this ratio.

Among rectangular tilings, we have explored two varieties: a regular grid, as shown in Figure 4c, and angled grids as shown in Figures 4d and 4e. Studies in perception show that the eye is least sensitive to grids angled at 45° and most sensitive to horizontal and vertical grids [14]. Thus, traditional screening methods have typically oriented the screen at an angle to reduce its visual impact. In the case of image mosaics, the image tiles are intended to be seen, and therefore the choice of whether or not to angle the grid is primarily an aesthetic decision.

4.3 Arranging the image tiles

Having selected a tiling grid, we need to place individual image tiles into the grid. We have explored a number of arrangement options:

- Use the same tile everywhere. (Figures 4c-e)
- Choose a random arrangement of the image tiles. (Figure 5)
- Arrange different tiles manually by eye. (Figure 2)
- Place image tiles by matching their average colors to the region of the target image that they cover. (Figure 8)
- Find a more detailed match between the tiles and the target image based on the shapes and colors within the images. (Figures 3 and 7)

For all but the first of these options, one has to consider whether a specific tile image may appear more than once in the final mosaic. Figure 7 does not permit repetition of the tile images; all other image mosaics shown here use repetition.

In order to convey the target image as effectively as possible, each location within the tile grid should contain the tile image that is most similar to the corresponding region of the target image. In this case, by “similar” we mean that the shapes, colors and textures of the tile image resemble those of the region in target image. Searching for the most similar tile image is a problem in content-based image retrieval, an area of active research. Several techniques that have been applied in this area are color histogram matching, texture analysis, shape analysis, edge matching, or a combination of these methods [10]. In creating Figures 3 and 7, we used the wavelets-based image matching algorithm due to Jacobs, Finkelstein and Salesin [6] to place tile images into the grid. The method uses wavelet analysis to distill each image down to a very small amount of data called a *signature*—essentially an extremely compressed version of the image, capturing only its broad forms and colors. Given the signature of a region in the target image, we can quickly search through the signatures of all possible tile images, choosing the one that best matches the target. To choose the best tile from among 20,000 images, the method requires less than a second on a conventional desktop workstation. It is beyond the scope of this paper to describe the method in detail; the reader is referred to [6] for a complete description.




Fig. 6. Shifting and scaling colors to darken an image tile. The distribution is shown as a histogram. Dashed line indicates average.

4.4 Color correction

Now that we have placed the image tiles in some arrangement, perhaps based on the shapes or colors of the target image, the next task is to alter their colors to better match the target image. In most of this section, we describe the correction process assuming that we are working with a grayscale image. At the end of this section we generalize the process for color images by using the same method in each of three color channels.

Our objective is to match the color (brightness, in the case of our grayscale image) of the tile image to the color of the region in the target image that is covered by the tile. If the target image has a constant color x across this region, then we want to adjust the color of the tile image so that its average color is x . If the brightness of the target image ranges from dark on the left side of the region to light on the right side of the region, then we would like the brightness of the image tile to somehow match this gradient as well. In addition to coarsely matching the colors of the target image, we would like to preserve as much as possible the features of the tile images. The approach we have taken is inspired by the dither matrix-based halftoning schemes described in Section 3.

Specifically, we make use of a *correction rule*, which takes as input an image tile and a desired average color a , and generates a *correction function* $F: R^1 \rightarrow R^1$ that maps a color x in the image tile to a color $F(x)$ in the final mosaic such that the region of the mosaic covered by the image tile will have the average color a . There are a variety of families of correction functions that could fulfill this role. For example, the constant function $F(x)=a$ would achieve the correct overall average, but ignores the original colors of the image tiles. A more sensible solution might be to scale the colors in the input tile so that the desired average is achieved. Suppose, for example, that the desired average color a is less than the average color a_t of the image tile. Then our “scaling” correction rule would generate the correction function $F(x)=(a/a_t)x$. A different correction rule—“shifting”—yields a correction function that shifts all the colors in the tile image as follows: $F(x)=x+(a-a_t)$. This scheme, shown schematically in the middle of Figure 6, will in general shift colors out of the range of reproducible colors. Other more sophisticated rules might employ gamma correction or a distortion of the color histogram [4], with improved effectiveness at the cost of additional complexity.

We have found that a combination of the “shifting” and “scaling” rules (shown schematically in Figure 6) is sufficient for our purposes and is easy to compute. Specifically, if we can use only a shift without sending any of the colors in the tile out of range, then we’re done. If not, we shift as much as possible, then scale the resulting

colors until the desired average is attained. So, let us return to the example where the desired average color a is less than the average color a_t of the image tile. Our “shift-and-scale” rule works as follows: if the minimum color m_t of the image tile is greater than $a_t - a$, then we use the shift rule above: $F(x) = x + (a - a_t)$; otherwise we use a combination of shifting and scaling: $F(x) = a(x - m_t) / (a_t - m_t)$. There is a symmetric pair of cases when the desired average a is greater than the tile average a_t .

Now that we have a correction rule, we simply apply that rule for every pixel in every tile in the mosaic, supplying as the desired average color a the color of the corresponding pixel in the target image. Observe that in regions where the target image has a fairly constant color, the tile image will simply be shifted and scaled to achieve the target color. In regions where more complex shade variations occur, the tile image will vary similarly. This scheme generalizes the traditional halftoning scheme described in Section 3. If we use a repeated pattern of the dither matrix D for our tiling, then we can generate traditional halftoning using the rule: $F(x) = \text{black}$ if $x < a$; $F(x) = \text{white}$ otherwise.

If the tile images are larger than 16x16 pixels, then an efficiency improvement may be implemented by building a table once for each tile image. The table contains the two parameters of the correction function—how much to shift, and how much to scale—for every possible desired average a . Then, for every pixel p in the tile image, we use the corresponding pixel in the target image as our index a into the table, and then shift and scale p according to the entries in the table.

So far we have described correction only for grayscale. Because we are working with images, we can perform the same kind of correction independently in each of three color channels, just as traditional halftoning is often performed in each of three (or four) color separations for printing. We have found that the resulting colors tend to be slightly more vibrant if the calculations are performed in the YIQ color space rather than the RGB color space [4]. As the results of using different color spaces is subtle at best, we have not investigated using other color spaces such as HSV or LAB [4]. Figures 5, 7 and 8 were produced by screening in each of the Y, I and Q color channels. A different variation was employed in Figure 3, in which the Y channel (which specifies the brightness of the image) was corrected as described above, while the I and Q channels (containing all of the *color* information) of the resulting mosaic were simply copied from the tile images. This tends to emphasize the tile images in the resulting mosaic.

5 Results

Figure 1 contains the only three-layer mosaic in this paper—an image of the face of Leonardo da Vinci’s *Mona Lisa* composed of 100 smaller images of her face. Each of the smaller images in turn is composed of 100 tiny images of the face. By implication, the layering in these images could be infinite, reminiscent of fractal structure [8]. We have not as yet experimented with *varied* imagery that is more than two layers deep.

US President John F. Kennedy appears in Figure 2 composed of images of Marilyn Monroe. This image also appeared in the 1994 Xerox PARC Algorithmic Art Show. In the creation of this image, final arrangement of the tile images was performed by hand, although an algorithmic matching based on the shapes in the image tiles proposed an initial arrangement.




Fig. 7. Blair Arch composed of 250 scenes at Princeton University. Target image is shown on the lower right. (Color image)

Grant Wood's painting *American Gothic* (1930) is composed in Figure 3 of pictures downloaded from the World Wide Web. The tile images were selected based on their colors and shapes from a collection of 20,000 images downloaded from all over the Web, using the image querying algorithm of Jacobs, *et al* [6].

In Figure 5, an abalone shell is shown screened by 16 microscopic images of abalone shells that were photographed using a scanning electron microscope. This figure demonstrates quite clearly the color correction process of Section 4.4, as the target image appears to burst through the tile images. In figures where the tiling is much more dense, this effect is more subtle.

In Figure 7, Princeton University's Blair Arch is shown composed of 250 scenes of the Princeton campus (in celebration of the University's 250th anniversary). The image tiles were arranged according to their shapes and colors as in Figure 3.

Finally, the globe shown in Figure 8 was created by texture-mapping [4] a sphere with a map of the earth. The map is actually an image mosaic composed of currency from around the world.



Fig. 8. Global currency (Color image)

The color correction process we describe is not computationally expensive: it is linear in the number of pixels in the image mosaic. Each of the figures in this paper was produced on a desktop workstation in a few minutes or less. Manually preparing the image tiles and the target image, as well as finding an arrangement for the tiles if it is done by hand tend to be the time-consuming stages in the overall process, particularly since they often are iterative.

6 Future work

This paper describes a method for creating image mosaics based on a target image and a collection of tile images. The project suggests a number of areas for future work, several of which are outlined here.

More complex tilings. As indicated in Sections 4.2 and 4.3, we use a fairly restrictive set of tilings: semiregular, rectangular grids. Perhaps we could find a more complex, irregular tiling that matches more carefully the tile images to a given target image. This

problem combines both continuous optimization (for the exact placement of the tile boundaries as well as the scaling and cropping of tile images) and discrete optimization (over the arrangement of tile images within the tiling).

3D image mosaic. We intend to construct a gallery-sized image mosaic installation in which the image tiles are posters distributed in a seemingly-haphazard arrangement at varying heights and depths in the environment. However, from a *single* vantage point, the posters will visually align to form the large mosaic.

Video mosaics. Image mosaics may be extended in the time dimension to create video mosaics. In these mosaics, the tiles may change over time, while the target image remains constant. Alternately, the target could evolve while the image tiles remain constant. One challenge for video mosaics is that the resolution of typical video is substantially lower than that of printed media; perhaps multiresolution video [3] could alleviate the resolution problem.

References

- [1] C. Close, J. Yau, *Recent Paintings*, Pace Wildenstein, 1995.
- [2] S. Dali, *The Salvador Dali Museum Collection*, Bulfinch Press, 1996.
- [3] A. Finkelstein, C.E. Jacobs, D.H. Salesin. Multiresolution Video. Proceedings of SIGGRAPH '96, *Computer Graphics*, Annual Conference Series, 1996, 281-290.
- [4] J. Foley, A. van Dam, S.K. Feiner, J.F. Hughes, *Computer Graphics, Principles and Practice, Second Edition*, Addison-Wesley, Reading, Massachusetts, 1990.
- [5] P.E. Haerberli. Paint by Numbers: Abstract Image Representations. Proceedings of SIGGRAPH '90, *Computer Graphics*, Annual Conference Series, 1990, 207-214.
- [6] C.E. Jacobs, A. Finkelstein, D.H. Salesin. Fast Multiresolution Image Querying. Proceedings of SIGGRAPH '95, *Computer Graphics*, Annual Conference Series, 1995, 277-286.
- [7] P. Litwinowicz. Processing Images and Video for an Impressionist Effect. Proceedings of SIGGRAPH '97, *Computer Graphics*, Annual Conference Series, 1997, 407-414.
- [8] B.B. Mandelbrot, *The Fractal Geometry of Nature*, W. H. Freeman, New York, 1983.
- [9] B. J. Meier. Painterly Rendering for Animation. Proceedings of SIGGRAPH '96, *Computer Graphics*, Annual Conference Series, 1996, 477-484.
- [10] W. Niblack, R. Barber, W. Equitz, M. Flickner, E. Glasman, D. Petkovic, P. Yanker, C. Faloutsos, G. Taubin. The QBIC Project: Querying Images By Content Using Color, Texture, and Shape. In *Storage and Retrieval for Image and Video Databases*, SPIE, 1993, 173-187.
- [11] V. Ostromoukhov, R.D. Hersch. Artistic Screening. Proceedings of SIGGRAPH '95, *Computer Graphics*, Annual Conference Series, 1995, 219-228.
- [12] S. Peleg, J. Herman. Panoramic Mosaics by Manifold Projection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1997, 338-343.
- [13] R. Silvers, M. Hawley, *Photomosaics*, Henry Holt and Co., 1997.
- [14] R. Ulichney, *Digital Halftoning*, MIT Press, Cambridge, Massachusetts, 1987.